

REMARKS

Claims 1-14 remain in the application. No claim has been allowed. Applicants thank the Examiner for withdrawing the objections to the claims and drawings.

Claims 1-14 were rejected under 35 U.S.C. 103(a) as being unpatentable over Baker, *et al.* (U.S. Patent 6,434,649) in view of Zwiegincew, *et al.* (U.S. Patent Application Publication 2005/0154705). It is believed that the Examiner failed to set forth a prima facie case of obviousness. Reconsideration is respectfully requested.

Claim 1 is directed to a Programmable Streaming Data Processor (PSDP) which is arranged to perform primitive functions directly on data received from a streaming data interface. The PSDP processes data from a streaming data source, such as a disk drive, prior to its being forwarded to a central processing unit (CPU) of a more general processor. The PSDP performs certain preliminary processing in order to reduce the computational load on the local CPU.

The PSDP may have processing logic known as a Data Engine that is capable of examining fields of a record to determine whether a record will or will not be passed to the CPU of the JPU as an output tuple. An output tuple is comprised of the fields of the source record from the disk that are to be selected for further processing by the CPU and PSDP generated fields. For example, a record retrieved from disk consists of a record header, typically containing more than one header field, and at least one data field, and typically, many data fields for each record. The collection of fields selected for return to the CPU as a result of processing a record is referred to a tuple. Possible tuple fields include various record header fields, the PSDP generated record address, unmodified record data fields, a hash field, and tuple status and length information. Boolean results and/or scratch pad words may also form parts of tuples. Most often a tuple will be shorter than the record that was used to generate it, but it may be longer, depending upon the program that is provided to the PSDP.

As data streams out of the filter, an output tuple is formed in a First In, First Out (FIFO) memory, in a way that permits aborting the tuple if the filter logic determines that the particular tuple should not be passed on to the CPU. Specifically, the memory FIFO has two write pointers, an “active” write pointer and a “visible” write pointer. The visible pointer maintains a position indicating a boundary of the last accepted tuple. Meanwhile, the active write pointer moves along the memory FIFO from the boundary, as words of the next possible tuple become available. As the PSDP logic determines that a tuple is not to be used, such as a result of the filter or TID processing described above, the memory FIFO’s active write pointer resets by moving back to the visible write pointer location. This has the effect of ignoring the intervening fields of the unwanted tuple and allowing them to be overwritten. If the PSDP logic makes this determination while the active pointer is still pointed to a field within the unwanted tuple, the active pointer will simply reset to the visible pointer location until the last field within that unwanted tuple has been overwritten.

If, on the other hand, the PSDP logic determines that a tuple is to be used, the visible pointer moves to the active pointer position, having the effect of keeping all intervening fields of the tuple that should be kept.

Baker, *et al.* relates to a data processor, and more specifically to a data transfer arrangement mechanism employed to transfer data to various components within a data processor. Such a multimedia processor and data transfer arrangement are directed to processing computer graphics and graphics on a standalone gaming console.

Zwiegincew, *et al.* relates to a method, system and computer program product for accessing data where the data store provider is abstracted from the data manipulation language used to access the data and the schema that describes the data. This prevents interdependencies such that changes in one technology do not require corresponding changes in another technology when combining multiple technologies.

However, there is no suggestion in Baker, *et al.* of a data engine, arranged to receive output data from the streaming interface FIFO, the data engine for determining field boundaries therein, and for processing fields to select one or more fields to be output tuples, the data engine also containing logic to determine whether an output tuple is to be selected for further processing by additional Job Processing Units (JPUs), and for asserting a use/lose decision value according to that determination.

More specifically, the claimed data engine, as a portion of the PSDP, provides programmable hardware directly in the disk read path from the controller and is programmed to understand the structure of the data that analysis software wishes to read and analyze. The PSDP can be programmed to operate on data as it is received from the disk, before it is stored in the JPU's memory, and in that process discard data that the CPU would otherwise have to analyze and discard in the absence of the data engine. In an embodiment specifically adapted for processing of field-oriented data from a database, data can be parsed by the PSDP into block header, record header, and record data fields of a database, and the record fields can be filtered by the PSDP so that only certain fields from certain records are actually forwarded to be written into the associated JPU's memory.

The data transfer switch of Baker, *et al.* is simply a routing mechanism for switching connections. All of the components of the multimedia processor in Baker, *et al.* are coupled to the data transfer switch: data cache, instruction cache, fixed function cache, memory controller, data streamer, PCI/AGP interface, and programmable input/output controller. The data transfer switch is not a data engine for determining field boundaries in output data, for processing fields to select one or more fields to be output tuples, does not contain logic to determine whether an output tuple is to be selected for further processing by additional JPUs, and is not for asserting a use/lose decision value according to that determination.

Thus, not only does the Applicant's invention require "determining field boundaries," but also requires "selecting fields as output tuples." Baker, *et al.* does not determine field boundaries or select fields to build output tuples. Baker, *et al.* might be "routing" data, but this is about all.

Further, Zwiegincew, *et al.* does not teach a tuple generator for assembling fields into the output tuple, and if the use/lose decision value indicates that such output tuple is to be discarded, from preventing such tuple set from being transferred from the output FIFO to the memory of the JPU. Zwiegincew, *et al.* teaches an operation performed on XML nodes specified in an XML data manipulation language. Never does Zwiegincew, *et al.* teach a tuple generator for assembling fields into the output tuple, let alone make use of tuples.

Moreover, there is no suggestion in Baker, *et al.* of an output FIFO device, for forming tuples and temporarily storing them prior to conditionally forwarding them to the JPU.

The Office Action is procedurally defective because the Examiner has failed to explain why one of ordinary skill in the art at the time the invention was made would have been motivated to combine the references.

The initial burden is on the examiner to provide some suggestion of the desirability of doing what the inventor has done. To support the conclusion that the claimed invention is directed to obvious subject matter, either the references must expressly or impliedly suggest the claimed invention or the examiner must present a convincing line of reasoning as to why the artisan would have found the claimed invention to have been obvious in light of the teachings of the references.

The Examiner nakedly asserts that one of ordinary skill in the art, at the time the invention was made to Zwiegincew, *et al.*'s method of manipulating data as a subcomponent to Baker, *et al.*'s Data Streamer, would have been motivated to combine Baker, *et al.* and Zwiegincew, *et al.* in order to improve data movement. The Examiner neither indicates how the references expressly or impliedly suggest the claimed invention nor presents a convincing line of reasoning as to why the artisan would have found the claimed invention to have been obvious in light of the teachings of the references.

The Examiner has failed to set forth a prima facie case of obviousness. There is no suggestion or motivation to combine reference teachings.

The suggestion or motivation cited by the Examiner to combine reference teachings does not lead to a teaching of Applicant's invention. Baker, *et al.* addresses the issue of a data movement arrangement that specifically accommodates data transfers for an integrated media processor chip set that contains various system components. However, Zwiegincew, *et al.* does not teach a method of improving data movement, as the Examiner claims. Rather, Zwiegincew, *et al.* teaches a method for accessing data where the data store provider is abstracted from the data manipulation language used to access the data and the schema that describes the data.

There is no express motivation to combine the references. The Examiner uses impermissible hindsight in examining the cited references. To reach a proper determination under 35 U.S.C. 103, the examiner must step backward in time and into the shoes worn by the hypothetical person of ordinary skill in the art when the invention was unknown and just before it was made. Knowledge of applicant's disclosure must be put aside in reaching this determination. The examiner must not use impermissible hindsight, developed by knowledge of the patent application, as support for patent claim rejections.

The Examiner has failed to set forth a prima facie case of obviousness. The prior art references do not teach or suggest all the claim limitations.

In order for the Examiner to cite combined prior art references in a prima facie case of obviousness, all claim limitations must be considered. By the Examiner's own admission, Baker *et al.* fails to explicitly teach the further limitation of the data engine wherein a use/lose decision value is utilized and the further limitation of the tuple generator assembling fields into the output tuple, and if the use/lose decision value indicates that such output tuple is to be discarded, for preventing such tuple set from being transferred.

The Examiner then points to Zwiegincew, *et al.* for a reference disclosing the further limitations of the data engine wherein a use/lose decision value is utilized and the further limitation of the tuple generator assembling fields into the output tuple, and if the use/lose decision value indicates that such output tuple is to be discarded, for preventing such tuple set from being transferred in order to improve data movement.

However, the combination of Baker, *et al.* and Zwiegincew, *et al.* fails to teach the further limitation of the use/lose decision value being asserted *according to that determination* (whether an output tuple is to be selected for further processing by additional processing Job Processing Units as determined by the data engine logic) and the further limitation of preventing such a tuple set from being transferred *from the output FIFO to the memory of the JPU* (if the use/lose decision value indicates that such output tuple from the tuple generator is to be discarded).

Rejections Under 35 U.S.C. §103(a)

Claims 1-14 were rejected under 35 U.S.C. 103(a) as being unpatentable over Baker, *et al.* (U.S. Patent 6,434,649) in view of Zwiegincew, *et al.* (U.S. Patent Application Publication 2005/0154705).

With regard to the Examiner's rejection of Claim 1, the combination of Baker, *et al.* with Zwiegincew, *et al.* does not overcome any of the deficiencies of Baker, *et al.* Further, the combination of Baker, *et al.* with Zwiegincew, *et al.* does not teach a data engine, arranged to receive output data from the streaming interface FIFO, the data engine for determining field boundaries therein, and for processing fields to select one or more fields to be output tuples, the data engine also containing logic to determine whether an output tuple is to be selected for further processing by additional Job Processing Units (JPUs), and for asserting a use/lose decision value according to that determination; a tuple generator for assembling fields into the output tuple, and if the use/lose decision value indicates that such output tuple is to be discarded, for preventing such tuple set from being transferred from the output FIFO to the memory of the

JPU; or an output FIFO device, for forming tuples and temporarily storing them prior to conditionally forwarding them to the JPU. For these reasons, the Examiner's rejection of Claim 1 is overcome and reconsideration is respectfully requested.

With regard to the Examiner's rejection of Claim 2, the combination of Baker, *et al.* with Zwiegincew, *et al.* does not overcome any of the deficiencies of Baker, *et al.* The combination of Baker, *et al.* with Zwiegincew, *et al.* fails to disclose use/lose values indicating a result from logic processing of fields read from the streaming data interface. Zwiegincew, *et al.* does not assign a use/lose value to data to determine whether the data engine should select it for an output tuple. Further, Claim 2 is dependent on Claim 1 and therefore contains all the limitations of the base claim. For these reasons, the Examiner's rejection of Claim 2 is overcome and reconsideration is respectfully requested.

With regard to the Examiner's rejection of Claim 3, the combination of Baker, *et al.* with Zwiegincew, *et al.* does not overcome any of the deficiencies of Baker, *et al.* The combination of Baker, *et al.* with Zwiegincew, *et al.* fails to disclose a Transaction Identifier. Further, Claim 3 is dependent on Claim 1 and therefore contains all the limitations of the base claim. For these reasons, the Examiner's rejection of Claim 3 is overcome and reconsideration is respectfully requested.

With regard to the Examiner's rejection of Claim 4, the combination of Baker, *et al.* with Zwiegincew, *et al.* does not overcome any of the deficiencies of Baker, *et al.* The combination of Baker, *et al.* with Zwiegincew, *et al.* fails to disclose TID processing and data engine logic being executed in parallel and makes no reference to TID processing or data engine logic. Further, Claim 4 is dependent on Claim 3, which is dependent on Claim 1, and therefore contains all the limitations of the base claim. For these reasons, the Examiner's rejection of Claim 4 is overcome and reconsideration is respectfully requested.

With regard to the Examiner's rejection of Claim 6, the combination of Baker, *et al.* with Zwiegincew, *et al.* does not overcome any of the deficiencies of Baker, *et al.* The combination

of Baker, *et al.* with Zwiegincew, *et al.* fails to disclose not asserting the use/lose decision value when a buffer local to the programmable data streaming processor is full; and means for appending an overflow filter bit to a tuple that indicates a transfer of a tuple that should be ignored. The value of the valid bit used in Baker, *et al.* indicates whether the specific byte is valid or not. Baker, *et al.* does not however, not assert a use/lose decision value when a buffer local to the programmable data streaming processor is full and make no use of an overflow filter bit to a tuple that indicates a transfer of a tuple that should be ignored. Baker, *et al.* makes no reference to tuples. Further, Claim 6 is dependent on Claim 5, which is dependent on Claim 1, and therefore contains all the limitations of the base claim. For these reasons, the Examiner's rejection of Claim 6 is overcome and reconsideration is respectfully requested.

With regard to the Examiner's rejection of Claim 9, the combination of Baker, *et al.* with Zwiegincew, *et al.* does not overcome any of the deficiencies of Baker, *et al.* The combination of Baker, *et al.* with Zwiegincew, *et al.* fails to disclose an overflow filter bit inserted in a length field appended to record fragments. Baker, *et al.* makes no reference to an overflow filter bit and does not insert an overflow filter bit in a length field appended to record fragments. Further, Claim 9 is dependent on Claim 1 and therefore contains all the limitations of the base claim. For these reasons, the Examiner's rejection of Claim 9 is overcome and reconsideration is respectfully requested.

With regard to the Examiner's rejection of Claim 10, the combination of Baker, *et al.* with Zwiegincew, *et al.* does not overcome any of the deficiencies of Baker, *et al.* The combination of Baker, *et al.* with Zwiegincew, *et al.* fails to disclose an invalid field appended to a tuple to indicate the results of transaction ID processing. Baker, *et al.* makes no reference to fields. Baker, *et al.* makes no reference to an invalid field, makes no reference to a tuple, and makes no reference to appending an invalid field to a tuple to indicate the results of transaction ID processing. Further, Claim 10 is dependent on Claim 1 and therefore contains all the limitations of the base claim. For these reasons, the Examiner's rejection of Claim 10 is overcome and reconsideration is respectfully requested.

With regard to the Examiner's rejections of Claims 5, 7-8 and 11-14, the combination of Baker, *et al.* with Zwiegincew, *et al.* does not overcome any of the deficiencies of Baker, *et al.* Further, Claims 5, 7-8 and 11-14 are either directly or indirectly dependent on Claim 1 and therefore contains all the limitations of the base claim. For these reasons, the Examiner's rejection of Claims 5, 7-8 and 11-14 is overcome and reconsideration is respectfully requested.

The combination of Baker, *et al.* with Zwiegincew, *et al.* does not overcome any of the deficiencies of Baker, *et al.* Therefore, the Examiner's rejection of Claims 1-14 is overcome and reconsideration is respectfully requested.

CONCLUSION

In view of the above amendments and remarks, it is believed that all claims are in condition for allowance, and it is respectfully requested that the application be passed to issue. If the Examiner feels that a telephone conference would expedite prosecution of this case, the Examiner is invited to call the undersigned.

Respectfully submitted,

HAMILTON, BROOK, SMITH & REYNOLDS, P.C.

By 

David J. Thibodeau, Jr.

Registration No. 31,671

Telephone: (978) 341-0036

Facsimile: (978) 341-0136

Concord, MA 01742-9133

Dated: 1/17/07